



Word frequency as a cue for identifying function words in infancy

Jean-Rémy Hochmann^{a,*}, Ansgar D. Endress^b, Jacques Mehler^a

^aSISSA, via Bonomea 265, 34136, Trieste TS, Italy

^bDepartment of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 43 Vassar St, 46-4127, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:

Received 20 July 2009

Revised 2 March 2010

Accepted 2 March 2010

Keywords:

Function words

Word frequency

Word learning

Syntactic bootstrapping

ABSTRACT

While content words (e.g., ‘dog’) tend to carry meaning, function words (e.g., ‘the’) mainly serve syntactic purposes. Here, we ask whether 17-month old infants can use one language-universal cue to identify function word candidates: their high frequency of occurrence. In Experiment 1, infants listened to a series of short, naturally recorded sentences in a foreign language (i.e., in French). In these sentences, two determiners appeared much more frequently than any content word. Following this, infants were presented with a visual object, and simultaneously with a word pair composed of a determiner and a noun. Results showed that infants associated the object more strongly with the infrequent noun than with the frequent determiner. That is, when presented with both the old object and a novel object, infants were more likely to orient towards the old object when hearing a label with a new determiner and the old noun compared to a label with a new noun and the old determiner. In Experiment 2, infants were tested using the same procedure as in Experiment 1, but without the initial exposure to French sentences. Under these conditions, infants did not preferentially associate the object with nouns, suggesting that the preferential association between nouns and objects does not result from specific acoustic or phonological properties. In line with various biases and heuristics involved in acquiring content words, we provide the first direct evidence that infants can use distributional cues, especially the high frequency of occurrence, to identify potential function words.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Language acquisition involves learning both syntax and the lexicon. While these components are often studied separately, they might be intimately linked during language acquisition, and might cross-fertilize each other. Function words (such as determiners and prepositions) illustrate this point particularly well. They clearly are words that have to be acquired. However, in contrast to content words (which must be linked to some semantic referent), function words mainly serve syntactic rather than semantic purposes. On the one hand, function words might thus impair word learning – because they are words that children

might try to learn, and yet they have no clear meaning they could be mapped onto. On the other hand, function words might facilitate word learning – by providing syntactic cues that might then be used for learning other (content) words. To use the syntactic cues associated with function words, however, infants need to *identify* them in the first place.

While different authors have uncovered different cues that tend to distinguish content and function words (Cutler, 1993; Shi, Morgan, & Allopenna, 1998), such proposals meet with two problems. First, to be useful for language acquisition, the cues must be available in any language a child might end up learning, and cannot be specific to a particular language (e.g., English). Second, early in life, infants need to be able to *use* such cues to identify function word candidates. Here, we start assessing these issues, asking whether infants can attribute different properties to

* Corresponding author.

E-mail address: hochmann@sissa.it (J.-R. Hochmann).

potential content words and function words based on a language-independent distributional property of function words, namely their high frequency of occurrence.

2. Words, syntax, the chicken and the egg

Children acquire both the syntax and the lexicon of their native language. However, different theories disagree on the relation between the development of syntax and that of the lexicon. Specifically, proponents of semantic bootstrapping (Pinker, 1984) and usage-based theories of language acquisition (Dabrowska, 2001; Tomasello, 2003) hold that vocabulary acquisition facilitates syntax learning, while proponents of syntactic bootstrapping accounts (Gillette, Gleitman, Gleitman, & Lederer, 1999; Gleitman, 1994; Gleitman & Gleitman, 1992; Landau & Gleitman, 1985) propose that syntax boosts vocabulary acquisition. We will now briefly review both kinds of theories.

Semantic bootstrapping theories describe how infants can bootstrap the initial steps of syntax acquisition based on their knowledge of (a limited number of) words. For example, semantic categories such as objects and actions might initially be used to discover how syntactic categories such as nouns and verbs are implemented in the language. Specifically, infants might first acquire a few words related to the objects and actions they observe. Then, they might use these words to learn the corresponding syntactic categories. For example, object names might be mapped onto nouns, and words describing actions onto verbs. Based on such a mapping, infants might discover crucial aspects of the syntactic organization of their native language. For instance, knowing the verb 'eat' and the noun 'cookie' might be sufficient to decide whether the object comes after the verb (e.g., "eat cookies", corresponding to the canonical English word order), or whether the object precedes the verb (e.g., "cookies eat", corresponding to the canonical Japanese word order; Pinker, 1994, p. 112). On this view, infants can start acquiring syntax only after having learned a minimal set of words, because knowledge of these words is crucially required to bootstrap grammar acquisition.

Semantic information might help grammar acquisition in yet another way. According to usage-based theories of syntax acquisition (see e.g., Dabrowska, 2001; Tomasello, 2003), infants and children first learn specific word sequences, with very limited knowledge of their underlying structure. That is, they might remember words only in specific contexts, and assign meaning to words only within this context. Crucially, however, as they do not analyze sentences in terms of their underlying structure, they should be unable to use words in contexts that differ from those they have heard. For example, if they have heard the word "broke" only in the sentence "The window broke", they should be unable to use the word in new contexts such as "He broke it" or "The windows got broken" (e.g., Savage, Lieven, Theakston, & Tomasello, 2003; Tomasello, 2003; but see Lidz, Gleitman, & Gleitman, 2003; Thothathiri & Snedeker, 2008). As they get older, children should gradually discover that the sentences they have heard have in fact an underlying structure, eventually leading to the kind of abstract syntactic knowledge observed in mature,

adult speakers. According to this theory, children thus need to acquire a substantial vocabulary before learning any syntactic regularity.

In contrast to such views, syntactic bootstrapping models hold that syntactic knowledge facilitates vocabulary acquisition (Gillette et al., 1999; Gleitman, 1994; Gleitman & Gleitman, 1992; Landau & Gleitman, 1985). For example, upon hearing a sentence like "the duck and the bunny are gorging", listeners as young as 2-year-olds are likely to conclude that "to gorp" must have an intransitive meaning, since it has no object. Upon hearing the sentence "the duck is gorging the bunny", in contrast, they tend to conclude that "to gorp" is transitive, since it now has a direct object (Naigles & Kako, 1993). Thus, a rather rudimentary syntactic analysis (such as counting the number of noun phrases and analyzing their positions) can constrain the interpretation of novel verbs.

Of course, semantic and syntactic bootstrapping accounts are not mutually exclusive, and infants might well use both routes in complementary ways. Both syntax and the lexicon might initially develop in parallel and cross-fertilize each other. This possibility is particularly important for the issue studied here, relating to how function words are acquired and used during language acquisition. From a syntactic bootstrapping perspective, the syntactic information carried by function words would be clearly helpful for learning new (content) words, as function words indicate syntactic roles and syntactic categories. For example, in a language like English, a word following a determiner is likely to be a noun, while a word following an auxiliary is likely to be a verb. Therefore, when hearing a novel word that is accompanied by a function word, infants might interpret it as referring to a novel object if the function word marks it as a noun (Brown, 1957), as referring to a novel action if the function word marks it as a verb (Bernal, Lidz, Millote, & Christophe, 2007; Brown, 1957), and as referring to a property when the function word marks it as an adjective (Waxman & Booth, 2001). This capacity seems to be present early in life, as infants as young as 14-month-old start using the syntactic information provided by function words to interpret new content words (Waxman & Booth, 2003).

While function words might facilitate the acquisition of content words by providing syntactic cues, they are more problematic from a vocabulary acquisition perspective, as they have no clear referents. As a result, unless infants can identify function words as function words, these words should impair vocabulary acquisition – because infants might consider them as meaningless "noise." In order to take advantage of the syntactic information provided by function words, infants thus need to identify them early on. In the next section, we will discuss a number of cues that might allow them to solve this problem.

3. Cues to identify function words

To identify function words, and to distinguish them from content words, infants might rely on two types of surface cues: phonological properties (Cutler, 1993; Shi, Werker, & Morgan, 1999; Shi et al., 1998), and distributional

cues (Gervain, Nespors, Mazuka, Horie, & Mehler, 2008; Shi et al., 1998).

Function and content words tend to have different phonological properties. Compared to content words, function words are often shorter, simpler and unstressed. In English, these differences are salient enough for neonates to notice them (Shi et al., 1999). However, even though such results demonstrate that very young infants are sensitive to these phonological differences, it remains unclear whether infants actually use them for language acquisition. In fact, the phonological differences between content and function words vary from one language to another. For example, English function words have reduced vowels (the [ðə], of [əv], etc.) and tend to start with certain consonants (th-, wh-) that are not commonly used in content words. In contrast, Hungarian function words do not have reduced vowels, and their initial consonants occur in many content words. In that language, function and content words differ mainly in the number of syllables (Gervain et al., 2008). In French, some function words are even homonymous with content words. For example, the sound [vo] can be both a determiner (as in “vos” – your) and a noun (as in “veau” – veal). Therefore, infants can use phonological cues to find function words only after having learned enough about their native language to identify the relevant cues.¹

Distributional cues, in contrast, seem to be relatively consistent across languages. Function words tend to occur at the edges of prosodic units, and therefore at utterance boundaries (Christophe, Millotte, Bernal, & Lidz, 2008; Gervain et al., 2008; Shi et al., 1998). Moreover, as all languages contain only a limited number of function words, their frequency of occurrence is much higher than that of content words. Indeed, except for proper names and nicknames such as ‘Mummy’ and ‘Daddy’, or non-referential interjection such as ‘look!’, the 20 or so most frequent words in child-directed speech are function words. This cue appears to be consistent across languages (see Gervain et al., 2008, for Italian and Japanese; Shi et al., 1998, for Mandarin Chinese and Turkish).

Gervain and colleagues (2008) showed that 7-month-old infants are sensitive to these distributional cues. In an artificial grammar learning experiment, they found that Italian infants preferred frequent elements to occur at the beginning of a unit (i.e., a bisyllabic word), whereas Japanese infants preferred the frequent elements to occur at the end of a unit. These preferences correlate with the word order of the participants’ native language. Indeed, (frequent) function words tend to occur at the beginning of units in Italian, especially in utterance-initial positions; in Japanese, in contrast, function words tend to occur at the end of units, especially in utterance-final positions. These results suggest that 7-month-old infants are sensitive to

variations of frequency of occurrence, and can use this cue to organize their input.

While these results suggest that distributional and, to a lesser extent, phonological cues to function words are available in infant-directed speech, and that, to some extent, infants seem to be able to process them, there is another question that has never been addressed: can infant learners actually use these cues to identify function words? This question is important, because there are numerous demonstrations showing that a perceptual sensitivity (such as the infants’ sensitivity to word frequency) is not necessarily used in all circumstances where it might be useful. For example, rats are sensitive to light flashes, as they can associate them with electroshocks; however, they cannot use this sensitivity to associate light flashes with visceral sickness. Conversely, they are sensitive to tastes, as they can associate them to visceral sickness; however, they cannot use this sensitivity to associate tastes with electroshocks (e.g., Garcia, Hankins, & Rusiniak, 1974). Hence, although rats are sensitive to both tastes and electroshocks, they cannot use these sensitivities for all kinds of associations.

The distinction between being sensitive to a cue and being able to use it is especially important for a cue such as frequency of occurrence. As many if not most animals are sensitive to this cue, this sensitivity likely evolved for non-linguistic reasons, raising the question of whether it can be used for aspects of syntax acquisition as well. For example, pigeons can categorize events by frequency (Keen & Machado, 1999; Machado & Cevik, 1997). Although they share the sensitivity to frequency of occurrence with human infants, they clearly cannot use it to acquire function words – because they do not acquire language. *Mutatis mutandis*, human infants might well be sensitive to acoustic or distributional differences between function words and content words – without using these potentially useful differences to discover function word candidates. Here, we start addressing this issue, asking whether a language-invariant cue to function words – their high frequency of occurrence – allows infants to attribute different properties to potential function words and to potential content words.

4. The current studies

Different cues have been proposed to be useful for identifying function words (Christophe et al., 2008; Gervain et al., 2008; Shi et al., 1998). However, it has never been shown whether infants can actually use them. Here, we start addressing this issue. We present infants with a word-learning situation, and ask whether they are more likely to attribute content-word like properties to infrequent items than to frequent items. Based on the hypothesis that it should be easier to associate objects with content words than with function word, we asked whether infants would be more likely to associate a visual object with a determiner or rather with a noun when listening to an unknown language.

In a pilot experiment (presented in Appendix A) we simply confirmed that the paradigm and stimuli used for

¹ Some prosodic cues might seem to be universal cues to function words. Specifically, function words are systematically less stressed than the content words they occur with (e.g., Nespors & Vogel, 1986). Note, however, that there are also some content words that receive systematically less stress than other content words; for example, in head-complement languages, pre-nominal adjectives tend to be less stressed than the nouns they occur with. It thus seems that watching out for “less stressed” words for identifying function words is not totally reliable.

and had no experience with French (the language used for our stimuli).

5.1.2. Stimuli

5.1.2.1. French sentences. In the familiarization phase, we presented infants with French sentences that each contained (in addition to a verb) two highly frequent determiners and two less frequent nouns. The list of words is presented in [Appendix B](#). Eighty-one sentences were created (see [Appendix C](#)). Each sentence was five words long and conformed to the following pattern: determiner–noun–verb–determiner–noun (e.g., “ce chat tue vos cerfs”; “this cat kills your deers”). The two determiners were always “ce” ([sə]; “this”) and “vos” ([vɔ]; “your”). Each determiner could be followed by nine different nouns, yielding to a total of 18 different determiner–noun combinations. Nine different verbs were used. [Appendix B](#) shows the phonetic transcriptions of all words, their number of occurrences in sentence-initial and sentence-final positions (that is, before or after the verb), respectively, and their English translations. All words were monosyllabic and had either a consonant–vowel or a consonant–vowel–consonant syllable structure. As a result, the determiners “ce” and “vos” were the most frequent words (and syllables) during the familiarization.

Forty-one sentences started with the determiner “ce” and 40 started with the determiner “vos”. The sentences were recorded from a female native speaker of French. The sentences, and all speech stimuli used in the present study, were recorded in a sound-attenuated chamber using a Sony ECM-S959C microphone connected to an M-Audio pre-amplifier and stored in the Audio Interface File Format (AIFF) (sample rate 44.1 kHz, sample size 16 bit). Files were normalized to a mean intensity of 60 dB using PRAAT (Version 5.0.25) sound processing software ([Boersma & Weenink, 2008](#)). Sentences lasted 1.4 s on average.

5.1.2.2. Objects. The objects used in the teaching and test phases were simple three-dimensional shapes generated as 3D animations in Maya 6.0 (Autodesk, Inc., San Rafael, CA), using a frame rate of 25 fps, the H.264 codec and the mov container format. [Fig. 2](#) shows the two objects we used. One was a blue three-dimensional cross. The other was a green pile of rings (similar to the belly of the emblem of the Michelin tire brand). Both objects were symmetrical, had similar perceived volumes and perimeters, and were found to be similarly attractive for 17-month old infants in a pilot study. However, the two objects had clearly different shapes and colors. During the teaching phase, one of the objects was presented in an animated movie of 33 s. During this movie, the object moved from one side to the other while rotating around its axes.

5.1.2.3. Labels. Four tokens of the phrases “ce chat” ([səʃa], “this cat”), “vos mets” ([vomɛ], “your dishes”), “ce met” ([səmɛ], “this dish”) and “vos chats” ([voʃa], “your cats”) were recorded from the same native speaker of French who also produced the familiarization sentences. (Note that in French, even though the spelling varies, the words “chat/chats” and “met/mets” are pronounced identically in the singular and plural forms.) The labels were recorded




Fig. 2. Objects used in all experiments. During the test phase, the two objects were presented statically as shown in the figure. During the teaching phase, in contrast, they slowly moved on the screen.

and stored in the same way as the French sentences described above. Labels had an average duration of 500 ms.

Two of these labels were used in the familiarization phase: “ce chat” or “vos mets”. The other two labels were used in the test phase: “ce met” and “vos chats”.

5.1.2.4. Apparatus. In this and all other experiments reported here, infants were tested individually. They sat on a parent’s lap 80 cm from a 17-in. LCD screen in a dimly lit, sound-attenuated cubicle. Parents wore dark sunglasses throughout the experiment to avoid all parental influence on the infants’ behavior.

5.1.3. Procedure

5.1.3.1. Familiarization. The 81 French sentences were played from a loudspeaker located behind the screen. They were presented in random order with a silence of 1 s between two sentences. While the sentences were played, the computer screen showed a silent movie. The movie consisted of a rotating chessboard changing its rotation direction every 30 s. The movie was chosen to be attractive enough to keep infants attentive, while containing no three-dimensional shape that could be interpreted as an object. The entire familiarization lasted 3 min 15 s.

5.1.3.2. Teaching phase. In the teaching phase, infants saw a visual object on the screen and simultaneously heard a speech label. We used the two shapes described above and two labels composed of a determiner and a noun (“ce chat” – [səʃa], “this cat”; “vos mets” – [vomɛ], “your dishes”). All four words had occurred during the familiarization phase. Moreover, during the familiarization phase, the nouns selected for the teaching phase had occurred only in the determiner–noun combinations that were also used in the teaching phase, but never with the other determiner. The four stimulus combinations resulting from two objects and two labels were counterbalanced across infants. Each infant learned only one object–label combination.

During the teaching phase, the object moved from one side to the other on an LCD screen while rotating around its axes. Simultaneously, a label was repeated 23 times. Two repetitions of the label were separated by about 900 ms of silence, yielding a total duration of 33 s.

The presentation movie could be interrupted if the infant looked away for more than 2 s. However, all the 16 infants included in the analysis looked at the entire movie without interruption.

5.1.3.3. Test phase. Following the familiarization phase and the teaching phase, infants completed four test trials. Two successive trials were separated by the presentation of the central fixation attractor (i.e., a white cross moving back and forth). The experimenter started each trial by pressing a key when the infant was looking at the central fixation attractor.

Once the experimenter started a trial, the fixation attractor disappeared, and the two objects presented in Fig. 2 appeared on the computer screen. The objects then remained static and visible on the screen for the entire trial duration. Two seconds after the trial started, the central fixation attractor appeared again for 3.32 s. A test label (see below) was first pronounced while the central fixation attractor was still visible on the screen; the offset of the label was synchronized with the attractor's disappearance. Following this, the test label was repeated five more times, two consecutive repetitions being separated by 1 s of silence.

Each of the objects on the screen occupied a surface of 8.5 cm × 9 cm. The centers of the two objects were separated by about 19 cm. For each infant, the position of objects was counterbalanced across trials. We also counterbalanced the position in which the blue object appears in the first trial across infants.

Each infants completed two test trials in the “same-determiner” condition, and two test trials in the “same-noun” condition. The two conditions were identical except for the test label used. In both conditions, the label differed only in one word from the label heard during the teaching phase. In the same-determiner condition, the label had the same determiner as during the teaching phase, but a different noun. In the same-noun condition, the label had the same noun as during the teaching phase, but a different determiner. Specifically, infants who had heard the label “ce chat” (“this cat”) during the teaching phase heard “ce met” (“this dish”) in the same-determiner condition, and “vos chats” (“your cats”) in the same-noun condition. Infant who had heard the label “vos mets” (“your dishes”) in the teaching phase, heard “vos chats” (“your cats”) in the same-determiner condition and “ce met” (“this dish”) in the same-noun condition. Importantly, while all words had occurred during the familiarization phase, the specific label phrases used during the test phase were all new, and had not occurred during the familiarization phase.

5.1.4. Analysis

We defined two windows of interest for the analysis of the infant eye gaze. Each was a square of 11 cm × 13 cm, centered on one of the objects. Infants' looking behavior was monitored using a Tobii 1720 Eye-tracker system

and the Clearview 2.5.1 software package. Only infants for whom the eye tracker data for each trial contained at least one fixation of at least 100 ms in one of the two windows of interest were included in the analysis. Infants not meeting these criteria were excluded for insufficient eye tracker data. Considering the four test trials, infants for whom more than 70% of the total fixation time was spent in only one of the windows of interest were considered to exhibit a side bias and were rejected from further analysis.

As dependent variable for the main analysis, we considered the first look to one of the objects following the first presentation of the test label. The first look was defined as the first uninterrupted fixation of at least 100 ms in one of the two windows of interest described above.

In each test trial, infants were coded either as first looking at the familiar object from the teaching phase, or as first looking at the novel object. Infants who did not look at any object in at least one of the four test trials were excluded from the analysis. For each infant and each condition, we computed the proportion of trials in which they first oriented towards the familiar object from the teaching phase. We then assessed whether these proportions differed between the conditions using a two-tailed Wilcoxon signed rank test.

We further analyzed the individual choices of each infant using a “distribution analysis”. Infants were categorized into three groups: (i) those associating the familiar object from the teaching phase more with the determiner than with the noun (“*determiner associators*”), (ii) those associating the familiar object more with the noun than with the determiner (“*noun associators*”), and (iii) those associating the familiar object equally with the noun and with the determiner (“*neutral associators*”). Determiner associators were infants who were more likely to look first towards the object from the teaching phase in the same-determiner condition than in the same-noun condition. Noun associators were infants who were more likely to look first towards the object from the teaching phase in the same-noun condition than in the same-determiner condition. Finally, neutral associators were infants who were equally likely to look first towards the object from the teaching phase in the same-noun condition and in the same-determiner condition.

We determined the expected distribution of noun associators, determiner associators and neutral associators, respectively, in the following way. In each trial, infants scored 1 if they first looked at the familiar object, and 0 if they first looked at the novel object. With two trials per condition, each infant could obtain a total score of 0, 1 or 2 in each condition (Note that all infants included in the analysis completed all four trials.) If infants oriented randomly towards the two objects in the test phase, the probabilities to orient first towards the familiar object on 0, 1 and 2 occasions were .25, .5 and .25, respectively for each condition (assuming that the conditions are statistically independent). To be a determiner associator, infants had to fall into one of the following three cases: (i) one initial orientation towards the familiar object in the same-determiner condition, and no such orientation in the same-noun condition; (ii) two initial orientations towards

the familiar object in the same-determiner condition, and one such orientation in the same-noun condition; (iii) two initial orientations towards the familiar object in the same-determiner condition, and no such orientations in the same-noun condition. Summing over the individual probabilities of these cases, the probability of being a determiner associator was $.5 * .25 + .25 * .5 + .25 * .25 = .3125$. Symmetrically, the probability of being a noun associator was also $.3125$. To be a neutral associator, infants had to orient to the familiar object as often in both conditions. These infants could orient twice to the familiar object for each condition, once for each condition, or zero times for each condition. The probability of being a neutral associator was, therefore, $.25 * .25 + .5 * .5 + .25 * .25 = .375$. Given these probabilities, we used an exact multinomial test to assess whether the observed distribution of determiner associators, noun associators and neutral associators was expected by chance.

5.2. Results

Fig. 3 shows how often infants first looked towards the familiar object from the teaching phase as a function of the experimental condition. Infants were significantly more likely to look first to the familiar object in the same-noun condition ($M = 1.38$, $SD = .72$, $Mdn = 1.5$) than in the same-determiner condition ($M = .69$, $SD = .70$, $Mdn = 1$), Wilcoxon signed rank test, $W = 12$, $p = 0.016$, $CI_{.95, Mdn}$ difference = 0, 1.

These results were further confirmed by the distribution analysis. Eleven infants out of 16 were noun associators, two were determiner associators, and three were




Fig. 3. Results of Experiment 1. Proportion of first looks directed towards the familiar object from the teaching phase. Dots represent the means of individual participants, diamonds sample averages, and the dotted line the chance level of 50%. When the label used during the test phase had the same noun but a different determiner as the label presented during the teaching phase, infants were more likely to fixate first the familiar object from the teaching phase compared to the condition where the label used during the test phase had the same determiner but a different noun as the label presented during the teaching phase.

neutral associators. This distribution was significantly different from the distribution expected by chance, $p = 0.0077$ (exact multinomial test).

To make sure that each infant's fixation pattern was not based on what they happened to fixate *before* the offset of the label, we also performed separate analyzes for trials where infants already fixated an object (rather than the central attractor) before the offset of the label. In total, there were 14 such trials (21.9% of the trials). Among these 14 trials, eight occurred in the same-determiner condition and six in the same-noun condition. After removing these 14 trials, our central pattern of results remained unchanged: infants were significantly more likely to look first to the familiar object in the same-noun condition ($M = 1.14$, $SD = .77$) than in the same-determiner condition ($M = .5$, $SD = .65$), Wilcoxon signed rank test, $W = 5$, $p = 0.020$.

5.3. Discussion

Experiment 1 asked whether infants could use cues provided by the speech signal to identify possible function words and possible content words, even when exposed to an unfamiliar language. Italian 17-month old infants were first exposed to a series of French sentences. Following this, they saw a visual object and heard simultaneously a bisyllabic label consisting of a highly frequent determiner and a less frequent noun (e.g., “ce chat”, “this cat”). We hypothesized that potential content words should be more likely to be associated with a referent than potential function words. If this is the case, and if Italian infants can use cues provided by French speech to identify function words, then they should associate the visual object more strongly with the noun from the label than with the determiner from the label.

To test this hypothesis, infants then saw two objects presented side-by-side. One was the visual object they had just seen, while the other one was novel. Simultaneously, infants heard a label derived from that used during the teaching phase. The label was derived either by changing the determiner and keeping the noun the same, or by changing the noun and keeping the determiner the same. Infants were more likely to first orient towards the familiar object when the derived label had the same noun as the label from the teaching phase, compared to the condition where the derived label had the same determiner as the label from the teaching phase. These results suggest that infants formed a stronger association between the object and the noun than between the object and the determiner.

While the results of Experiment 1 suggest that infants treated words with a relatively low frequency as more content-word like than words with a relatively high frequency, it is less clear which cues infants used to identify potential function words. Infants might have used the frequency of the determiners to identify them as potential function words, as the determiners were much more frequent than any other content word in the familiarization phase of Experiment 1. However, they might also have used a different cue. Specifically, in French, as in most languages of the world, content words are usually more salient than

function words because they bear the main prominence (Nespor & Vogel, 1986). If infants pay more attention to more salient syllables than to less salient syllables, they might have associated the object more strongly with nouns than with determiners for this reason, and not because of their sensitivity to distributional information.

Experiment 2 was designed to control for this possibility. Experiment 2 was identical to Experiment 1 except that infants did not undergo the familiarization phase, but only the teaching and test phases. Crucially, these two phases used exactly the same stimuli as in Experiment 1. As the stimuli used in these phases were naturally recorded, the nouns were still more salient than the determiners. Hence, if acoustical properties such as stress, or any other property specific to the stimuli of the teaching and test phases, are sufficient to account for the infants' preferential association between nouns and the visual objects, we would expect similar results as in Experiment 1. Conversely, if the distributional properties such as the frequency of occurrence determined infants' behavior in Experiment 1, we would expect different results in Experiment 2.

6. Experiment 2

6.1. Materials and methods

Experiment 2 was identical to Experiment 1, except that there was no familiarization phase prior to the teaching and test phases.

Twenty-one new Italian 17-month old infants were tested. Five were excluded for fussiness (1), equipment failure (1) or side bias (3). The remaining 16 infants (5 males, 11 females, age range: 17 months and 2 days – 17 months and 22 days) were included into the final analysis. As in Experiment 1, these infants acquired Italian and had no experience with French (the language we used for our stimuli).

6.2. Results

Fig. 4 shows how often infants first looked towards the familiar object from the teaching phase as a function of the experimental condition. Infants started looking towards the familiar object from the teaching phase as often in the same-noun condition ($M = 1.13$ $SD = .62$, $Mdn = 1$) as in the same-determiner condition ($M = 1.06$, $SD = .68$, $Mdn = 1$), Wilcoxon signed rank test, $W = 25$, $p = 0.82$, $CI_{.95, Mdn}$ difference = $-.5, .5$.

We further analyzed the results of Experiment 2 using the distribution analysis. Six infants out of 16 were noun associators, five were determiner associators, and five were neutral associators. This distribution was not significantly different from chance, as assessed by an exact multinomial test, $p > 0.9$.

Fig. 5 shows the comparison of Experiments 1 and 2. We calculated for each infant the difference in proportions of first orienting towards the familiar object in the same-noun condition and in the same-determiner condition, respectively. This difference in Experiment 1 was marginally different from that in Experiment 2, Wilcoxon rank sum test, $U = 313$, $p = 0.054$, $CI_{.95, Mdn}$ difference = $0, 1$.

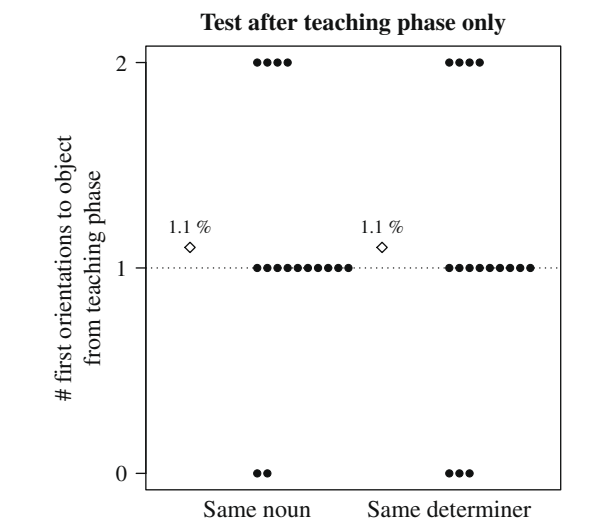


Fig. 4. Results of Experiment 2. Proportion of first looks directed towards the familiar object from the teaching phase. Dots represent the means of individual participants, diamonds sample averages, and the dotted line the chance level of 50%. During the test phase, infants were as likely to fixate first the familiar object from the teaching phase as to fixate first a new object, both when the label used during the test phase had the same noun but a different determiner as the label presented during the teaching phase, and when the label used during the test phase had the same determiner but a different noun as the label presented during the teaching phase.




Fig. 5. Results of Experiments 1 and 2. Proportion of first looks directed towards the familiar object from the teaching phase. Bars show the sample averages in the different conditions and errors bars the standard errors from the average.

ally different from that in Experiment 2, Wilcoxon rank sum test, $U = 313$, $p = 0.054$, $CI_{.95, Mdn}$ difference = $0, 1$.

To make sure that each infant's fixation pattern was not based on what they happened to fixate before the offset of the label, we performed separate analyzes for trials where infants already fixated an object (rather than the central attractor) before the offset of the label. Infants were not looking at the central fixation when the test label was

pronounced in seven trials (10.94% of all trials). Among these seven trials, four occurred in the same-determiner condition and three in the same-noun condition. After removing these seven trials, our central pattern of results remained unchanged: infants started looking towards the familiar object from the teaching phase as often in the same-noun condition ($M = 1.063$, $SD = .57$) as in the same-determiner condition ($M = .878$, $SD = .72$), Wilcoxon signed rank test, $W = 12$, $p = 0.562$. In that analysis, the difference between Experiments 1 and 2 only showed a trend, Wilcoxon rank sum test, $U = 254$, $p = 0.104$.

6.3. Discussion

In Experiment 1, infants preferentially associated visual objects with nouns rather than with determiners. A possible conclusion from this experiment is that infants might have used the very high frequency of the determiners to identify them as function words, and thus to disprefer them as carriers of meaning. However, infants might also have relied on acoustic or phonological properties of the words used in the teaching and test phases, especially on the greater saliency of nouns compared to determiners.

Experiment 2 controlled for this possibility by replicating Experiment 1 without the initial familiarization. As a result, Experiment 2 provided infants with all acoustic and phonological cues present in Experiment 1, but determiners were no longer more frequent than nouns.

If the results of Experiment 1 had been carried exclusively by the greater acoustic saliency of nouns relative to determiners, one would expect similar results in Experiments 1 and 2. In contrast, if distributional information, and especially the high frequency of determiners, influenced the infants' behavior in Experiment 1, the results in Experiment 2 should differ from those in Experiment 1.

In contrast to Experiment 1, where infants' first fixations tended to be directed towards the familiar object from the teaching phase in the same-noun condition, and towards the novel object in the same-determiner condition, infants' first fixations in Experiment 2 were directed equally often to either object. It thus seems safe to conclude that, in Experiment 2, infants did not preferentially associate visual objects with nouns compared to determiners. Hence, the high frequency of determiners in the familiarization phase of Experiment 1 seems to play a crucial role for establishing the preferential association between visual objects and nouns.

This interpretation is strengthened when comparing the results of Experiment 2 to those from the pilot experiment reported in [Appendix A](#). The pilot experiment was similar to Experiment 2, except that, during the test phase, infants were presented either with a label that was identical to that used during the teaching phase, or with an entirely novel label with both a different determiner and a different noun (as opposed to Experiments 1 and 2, where the test labels changed either the determiner or the noun, but not both). Compared to the condition with the old label from the teaching phase, the infants' first look was directed more often towards the novel object when presented with the entirely novel label. These results contrast with those of Experiment 2, where infants were equally likely to ori-

ent to either object in both conditions. A tentative conclusion is that infants recognized that parts of the test labels in Experiment 2 were identical to the labels used during the teaching phase; in fact, if they had considered a label with a changed noun or with a changed determiner as a novel label, they should orient towards the novel object as in the pilot experiment.

Hence, although infants in Experiment 2 seemed to recognize that parts of the test labels were identical to the label used during the teaching phase, and although all acoustic and phonological cues provided in Experiment 1 were also present in Experiment 2, we did not observe the preferential association between visual objects and nouns as in Experiment 1. Together, these results thus suggest that this preferential association cannot be explained based exclusively on acoustic or phonological properties of the labels. Rather, the combined results of Experiments 1 and 2 suggest that the distributional cues implemented in the familiarization phase of Experiment 1 influenced the infants' behavior in the successive phases. Specifically, we suggest that infants prefer to associate a novel object with an infrequent word rather than with a frequent word, because infrequent words are more likely to be content words and frequent words are more likely to be function words.

7. General discussion

The most frequent words in a language (i.e., determiners, prepositions, pronouns and conjunctions) do not refer to any particular perceivable objects, actions or properties. Rather, they are mainly in the service of syntax. In line with their largely syntactic role, these words are called function words, as opposed to content words such as nouns, verbs, adjectives and adverbs. Function words are characterized by various distributional and phonological properties. However, it is unknown whether young infants can actually use such properties to identify potential function words. In the present study, we started addressing this issue, asking whether 17-month old infants can identify potential function words, based on a language-universal cue, namely their high frequency of occurrence.

In Experiment 1, we presented infants with short sentences in a language unknown to them. These sentences contained frequent and infrequent words, corresponding to actual function and content words, respectively. We asked whether infants were more likely to associate the relatively infrequent words with external referents as compared to the relatively frequent words. If infants consider highly frequent words as potential function words, they might be less likely to associate them with objects – because function words are less likely than content words to have referents in natural languages. Moreover, if they consider relatively infrequent words as potential content words, they might be more likely to associate them with objects. Our results suggest that infants associated the object more strongly with the infrequent noun than with the frequent determiner.

When the initial familiarization to French sentences was removed, however, infants no longer preferentially

associated the object with the noun, although the results of a pilot experiment (reported in [Appendix A](#)) shows that infants can associate labels with objects in the absence of familiarization. Hence, the combined results of Experiment 2 and the pilot experiment show that the preferential association between content words and objects observed in Experiment 1 cannot be explained solely due to phonetic or phonological properties of the function words and content words. Rather, our results suggest that infants can use the distributional properties of words in their input to infer certain of their properties. Infants are less likely to associate objects with very frequent words than with relatively less frequent words, thus raising the possibility that they might consider these words as function word candidates.² Below, we will discuss possible mechanisms that might account for these results, and that might also contribute to the identification of function words in natural language acquisition.

7.1. Word frequency vs. other distributional cues

While our results seem to suggest that infants attribute function-word like properties to very frequent words, there might be other distributional cues that infants might exploit as well. Specifically, Italian infants (and adults) prefer frequent items to occur in unit-initial positions ([Gervain et al., 2008](#)), even though it remains unclear whether infants just learned a frequency pattern from their native language, or whether they linked this distributional regularity to the fact that function words tend to occur in unit-initial positions in Italian. If infants linked the frequency pattern to the distribution of function words, they might use this information in our experiments as well: they might have learned that the initial word of a phrase is generally not referential.

At first sight, Experiment 2 seems to rule out this possibility; after all, if infants simply considered the first part of a label as non-referential, they should do so in Experiment 2 as well as in Experiment 1. However, infants might require the familiarization phase of Experiment 1 to understand that labels are composed of two words – rather than of a single, bisyllabic word. Specifically, they might notice that the transitional probabilities between determiners and nouns are low; according to many authors (e.g., [Aslin, Saffran, & Newport, 1998](#); [Saffran, Aslin, & Newport, 1996](#)), they might postulate word boundaries where transitional probabilities are low. As a result, infants might consider the first part of the labels as non-referential in Experiment 1 (where it would correspond to a word), but not in Experiment 2 (where it would correspond to the first syllable of a bisyllabic word).

While our results do not allow us to rule out this possibility, we believe that it is unlikely for at least two reasons.

² Another way to explain our results is the possibility that infants might consider relatively infrequent words as better candidate for content words. However, infants participating in Experiment 1 are more likely to have extracted the frequent, rather than the infrequent words from the familiarization phase. Therefore, we consider that if they attributed a property to some words, they must have attributed a property to the frequent words.

First, it is debated whether learners really postulate word boundaries at locations of low transitional probabilities (see [Endress & Mehler, 2009](#), for experiments with speech stimuli; see [Vickery & Jiang, 2009](#), for related results with visual stimuli; see [Yang, 2004](#), for computational arguments). If infants cannot use transitional probabilities to postulate word boundaries, transitional probabilities cannot help them either to understand that labels are composed of two words.

Secondly, since the familiarization in Experiment 1 was composed of naturally recorded sentences, determiners and nouns formed perceptual groups due to the prosodic organization of the sentences (e.g., [Nespor & Vogel, 1986](#)). Even if infants can use transitional probabilities to postulate word boundaries, prosodic cues dominated distributional cues such as transitional probabilities in all experiments that have tested this issue (e.g., [Johnson & Jusczyk, 2001](#); [Johnson & Seidl, 2009](#); [Shukla, Nespor, & Mehler, 2007](#); [Thiessen & Saffran, 2003](#)). As a result, it is not clear whether the familiarization in Experiment 1 should make it harder to split the labels into two words (because determiners and nouns occur in prosodic groups), or whether it should make it easier (because transitional probabilities between determiners and nouns are low).

While further research is needed to determine whether infants can use transitional probabilities to split labels into two words, and whether this capacity explains the differences between Experiments 1 and 2, our results provide the first direct demonstration that infants can use distributional information to attribute function-word like properties to specific elements in their speech input, raising the possibility that they might follow similar strategies in natural language acquisition as well.

7.2. Word frequency influence on word–object associations

Why are 17-month old infants more likely to associate an object with an infrequent word rather than with a frequent one? Below, we will consider three possible reasons and conclude that, despite being at first sight plausible accounts of our data, neither is fully satisfactory. We then suggest that, just as children use various biases and heuristics for learning content words ([Markman, 1990](#); [Hirsh-Pasek, Golinkoff, & Hollich, 2000](#)), they might also have universal biases for identifying function words, especially a bias to consider highly frequent words as function word candidates.

7.2.1. Can “mutual exclusivity” account for the preferential association of objects with nouns?

A first possible explanation of our results is related to an observation known as the mutual exclusivity assumption ([Markman, 1990](#)). When learning words, infants assume that one type of objects has only one label ([Halberda, 2003](#); [Markman, 1990](#)). [Liittschwager and Markman \(1994\)](#), for instance, showed that 16-month-olds find it more difficult to learn a new label for an object they can already name, compared to learning a new label for an unfamiliar object. As a result, when infants have to find the referent of a novel word, they tend to choose an object for which they have no previous name, rather than an al-

ready labeled object (Markman, Wasow, & Hansen, 2003). Potentially, the results of Experiment 1 might be due to a very similar principle: in addition to assuming that one type of object has only one name, they might also assume that one label refers only to one type of object. If so, infants might choose a label that is not yet associated with a referent when choosing among possible labels for a novel object.

At first glance, this possibility seems to provide a plausible account of the results of Experiment 1. Indeed, as function words occurred much more often than content words during the familiarization to French sentences, infants had many more opportunities to associate them with some referent. Once they had learned a referent for a function word, infants might not be willing to associate this word to a second referent. Therefore, when a novel object was accompanied by a label composed of a determiner and a noun, and if the determiner was already associated to some other referent, infants should have been more likely to associate the novel object with the noun rather than the determiner.

While plausible, the explanation can hold only if infants had already assigned a meaning to the determiners *before* the teaching phase of Experiment 1. This, however, is unlikely to be the case for several reasons. First, our participants had never heard French before taking part in the experiment, and none of the words we used have a meaning in their native language (i.e., Italian). Therefore, the infants could have learned a meaning for one or both of the determiners only during the familiarization phase. During the familiarization phase, however, the infants' attention was attracted mainly by a rotating chessboard-like display presented on the screen. This bi-dimensional display, however, was chosen because it is, we believe, unlikely to be perceived as an object (see also Stager & Werker, 1997). Given that the display was unlikely to be perceived as an object, it is also unlikely that the infants associated it with one of the determiners.

Second, the determiners were never presented in isolation; rather they were embedded in sentences and surrounded by other French words, each of the determiners representing only 20% of the words heard during the familiarization. To the best of our knowledge, there is no evidence that 17-month-old infants can associate a word and a visual stimulus in such a situation. Therefore, participants are unlikely to have associated the frequently repeated determiners with the visual attractor or to any object. Hence, the mutual exclusivity principle does not seem to provide an explanation for our results.

7.2.2. Can “cross-situation statistics” account for the preferential association of objects with nouns?

A second possible explanation of our results is related to what has been called cross-situational statistical learning. For learning the meaning of words, children might track the frequency with which a word and an object co-occur (Smith & Yu, 2007; Voloumanos, 2008; Yu & Smith, 2007). In the most commonly assumed form of cross-situational statistics, learners keep track of the number of times a given word and a given object occur together. According to the axiom that what fires together wires to-

gether, the associative link between the word and the object representations is strengthened each time they occur together. The meaning of a word is then determined by selecting the referent with the strongest associative strength in a “winner-takes-all” fashion.

This simple view of cross-situational statistics does not account for our data. Indeed, the teaching phase is the only phase where infants were likely to form associations between objects and their labels. In that phase, the object occurred as often with the determiner as with the noun. According to the version of the cross-situational statistics outlined above, the associations between the determiner and the object should thus be as strong as between the noun and the object. Therefore this view does not provide an account for our data.

There is an alternative version of cross-situational statistics, however, that seems more consistent with our results. In fact, the aforementioned version of cross-situational statistics does not take into account how often the object appears without any word being pronounced, and how often a word is uttered in the absence of any object. Alternatively, the associative strength between an object and its label might be determined by counting how often the object and the label co-occur, and by comparing this number to the number of times the object and the label are encountered in total. (Technically, such a measure of associative strength could be expressed in terms of mutual information or related statistics.) The label of an object might then still be determined by selecting the label with the strongest associative link.

This version of cross-situational statistics seems to be consistent with our data. Indeed, recall that each determiner occurred nine times as often as each noun during the familiarization phase. During the teaching phase, in contrast, only one determiner and one noun occurred, and they occurred equally often. Moreover, each determiner appeared 81 times in the familiarization and only 24 times in the teaching phase. Compared to the noun, the determiner thus appeared much more often in the absence of a referent than with a referent, which in turn should weaken the associations it can form with any referent. Accordingly, the association of the object with the noun should therefore be stronger than the association with the determiner.

While this account is in agreement with our results, it is inconsistent with previous data. Specifically, this account predicts that frequently hearing a word in absence of its referent should *always* be detrimental for acquiring its meaning, a prediction that seems problematic both based on every day experience, and on previous experiments. For example, Swingley (2007) familiarized 18-month-old infants with a novel word form, but, crucially, without presenting its referent. Following this familiarization, he administered a word learning task where infants had the opportunity to associate a word with visually a presented item. Results showed that infants were better at learning a referent for the word they had been familiarized with compared to learning a referent for a novel word. In contrast to the predictions of the cross-situational statistics model, hearing a word in the absence of its referent is thus not necessarily detrimental for associating this word with

a meaning, as infants were *better* at learning the meanings for the words that they had heard frequently.

Considering Swingley's (2007) data together with our own, we suggest the following account of how infants might use frequency as a cue to function words. A first step to learning words is presumably to acquire their phonological representations, as infants need some phonological representation to map it to a meaning. Swingley's (2007) data suggest that hearing a word frequently is conducive to acquiring its meaning, presumably because hearing it frequently allows learners to consolidate its phonological representation. Once such a phonological representation is in place, however, frequency might play a different role for acquiring the word's meaning. Presumably, children try (explicitly or implicitly) to guess the meanings of words as they hear them. As function words occur very frequently with no referent that is consistently associated with them, children should have difficulties figuring out their meaning, and, at some point, might simply "give up" looking for meanings of function words. That is, due to their high frequency, children should remember the phonological form of function words fairly well. However, once children have heard these words for some critical number of times without being able to figure out their meaning, they might tag their phonological forms as uninformative or non-referential and, perhaps, as potential function words.

7.3. Conclusion

In conclusion, the present study shows that, when exposed to an unknown language, infants attribute to highly frequent words at least one characteristic of function words: they are not associated to specific referents such as objects. Just as infants use different biases to learn content words, they might use another bias to identify function words that exploits one of their language-universal properties: their high frequency. Upon frequently encountering words that have no clear meaning, infants might interpret such words as non-referential function-word candidates. This bias might represent the initial step to discriminating function words from content words. This might in turn allow infants to focus on the content words when enriching their vocabulary, while focusing on function words to build syntactic relations.

Acknowledgement

This research has been supported by McDonnell Foundation Grant No. 21002089, the European Commission Special Targeted Project CALACEI (Contract No. 12778 NEST) and the Mind, Brain and Behavior Initiative at Harvard University. We thank Liuba Papeo for comments on a previous version of the manuscript, Alessio Isaja for technical support, and Francesca Gandolfo and Marijana Sjekloća for their help in testing infants.

Appendix A. Pilot experiment

The goal of this pilot experiment was to ensure that the parameters used in the teaching and test phases of Exper-

iments 1 and 2 would allow infants to learn the association between objects and labels. The experiment was identical to Experiment 2, except that the test labels were either identical to those used in the teaching phase, or fully different from those used in the teaching phase. That is, in the same-label condition, the test labels comprised both the same determiner and the same noun as the label from the teaching phase; in the novel-label condition, in contrast, both the determiner and the noun were different.

A.1. Materials and methods

A.1.1. Participants

Twenty-five Italian 17-month old infants were tested. Nine were excluded for fussiness (7) or equipment failure (2). The remaining 16 infants (6 males, 10 females, age range: 17 months and 02 days – 17 months and 27 days) were included in the final analysis. Importantly, these infants acquired Italian and had no experience with French (the language we used for our stimuli).

A.1.2. Stimuli

The same labels and objects as in Experiments 1 and 2 were used.

A.1.3. Procedure

Teaching phase: The procedure for this phase was identical to that in Experiments 1 and 2.

Test phase: There were two test conditions: the same-label condition and the novel-label condition. The two conditions were identical except for the test label used. As mentioned above, in the same-label condition, the label




Fig. A1. Results of the pilot experiment presented in Appendix A. Proportion of first looks directed towards the familiar object from the teaching phase. Dots represent the means of individual participants, diamonds sample averages, and the dotted line the chance level of 50%. When the label used during the test phase was identical to that presented during the teaching phase, infants were more likely to fixate first the familiar object from the teaching phase compared to the condition where the label used during the test phase differed from that presented during the teaching phase.

was identical to the one used in the teaching phase. In the novel-label condition, the label was composed of a different determiner *and* a different noun. Specifically, infants who heard the label “ce chat” (“this cat”) during the teaching phase heard “ce chat” (“this cat”) in the same-label condition, and “vos mets” (“your dishes”) in the novel-label condition. Infant who heard the label “vos mets” (“your dishes”) in the teaching phase, heard “vos mets” (“your dishes”) in the same-label condition, and “ce chat” (“this cat”) in the novel-label condition. Infants completed two trials in the same-label condition, and two trials in the novel-label condition.

Except for the labels, the procedure of the test phase was identical to that employed in Experiments 1 and 2.

A.2. Results

There was no trial in which infants already fixated an object (rather than the central attractor) before the offset of the label. Fig. A1 shows how often infants first looked towards the familiar object from the teaching phase as a function of the experimental condition. Infants were significantly more likely to first look at the object from the teaching phase in the same-label condition ($M = 1.19$,

$SD = .66$, $Mdn = 1$) than in the novel-label condition ($M = .62$, $SD = .62$, $Mdn = 1$). Wilcoxon signed rank test, $W = 50$, $p = 0.015$, $CI_{.95, Mdn} \text{ difference} = 0, 1$.

These results were further confirmed by the distribution analysis. Nine infants out of 16 directed their first fixation more often towards the object from the teaching phase in the same-label condition than in the novel-label condition, one infant directed the first fixation more towards that object in the novel-label condition than in the same-label condition, and six infants directed their first fixation equally often towards either object in both conditions. This distribution was significantly different from chance, $p = 0.025$ (exact multinomial test).

Appendix B. List of words

Table B1.

Appendix C. List of sentences

Table C1.

Table C1

List of sentences used in the familiarization phase of Experiment 1.

Table B1

Words used in the sentences of the familiarization phase of Experiment 1.

Word	Pronunciation	Occurrences in sentences beginning with “ce”	Occurrences in sentences beginning with “vos”	English translation
<i>Nouns following the determiner “ce”</i>				
Rat	(ʁa)	5	4	Rat
Chat	(ʃa)	4	5	Cat
Bas	(ba)	5	4	Stocking
Pot	(po)	5	4	Jar
Show	(ʃo)	5	4	Show
Corps	(kɔʁ)	4	4	Body
Pet	(pɛ)	4	5	Fart
Fait	(fɛ)	4	5	Fact
Verre	(vɛʁ)	5	5	Glass
<i>Nouns following the determiner “vos”</i>				
Cas	(ka)	5	5	Cases
Pas	(pa)	5	4	Steps
Mas	(ma)	4	4	Provençal houses
Mots	(mo)	4	5	Words
Seaux	(so)	5	4	Buckets
Bords	(bɔʁ)	5	5	Sides
Mets	(mɛ)	5	4	Dishes
Quais	(kɛ)	4	4	Platforms
Cerfs	(sɛʁ)	4	5	Stags
<i>Verbs</i>				
Lie/lient	(li)	3	5	Links/link
Scie/scient	(si)	4	5	Saws/saw
Nie/nient	(ni)	5	5	Denies/deny
Noue/nouent	(nu)	5	3	Knot/knot
Joue/jouent	(ʒu)	4	5	Plays/play
Loue/louent	(lu)	5	4	Praises/praise
Voit/voient	(vwa)	5	5	Sees/see
Noit/noient	(nwa)	4	4	Drowns/drown
Tue/tuent	(ty)	6	4	Kills/kill

Sentences beginning with “ce”	Sentences beginning with “vos”
Ce rat tue vos cerfs	Vos mas lient ce rat
Ce pet joue vos quais	Vos quais voient ce rat
Ce show noie vos bords	Vos seaux noient ce show
Ce verre tue vos seaux	Vos cas tuent ce rat
Ce fait lie vos seaux	Vos bords scient ce chat
Ce pot scie vos pas	Vos mas louent ce corps
Ce pot voit vos bords	Vos pas voient ce fait
Ce pet noie vos mas	Vos mots noient ce pet
Ce bas tue vos cas	Vos cas noient ce bas
Ce rat voit vos mets	Vos cas scient ce corps
Ce fait loue vos quais	Vos pas louent ce verre
Ce corps tue vos bords	Vos mets scient ce fait
Ce pet nie vos seaux	Vos quais nient ce corps
Ce corps nie vos mets	Vos mots lient ce bas
Ce corps voit vos mots	Vos mets nient ce verre
Ce corps loue vos pas	Vos mas jouent ce show
Ce bas nie vos pas	Vos bords jouent ce bas
Ce bas noue vos mets	Vos cerfs scient ce verre
Ce show noue vos mots	Vos mets jouent ce chat
Ce chat noue vos quais	Vos bords nient ce pet
Ce pot nie vos cerfs	Vos mots tuent ce fait
Ce corp scie vos cerfs	Vos quais noient ce pot
Ce fait voit vos cas	Vos bords tuent ce verre
Ce show lie vos cerfs	Vos mas scient ce pot
Ce pet noue vos cas	Vos cerfs voient ce bas
Ce bas joue vos seaux	Vos bords lient ce fait
Ce chat nie vos mas	Vos mots voient ce verre
Ce chat noie vos cas	Vos pas noient ce chat
Ce chat scie vos seaux	Vos cerfs louent ce fait
Ce show joue vos pas	Vos mots nient ce chat
Ce fait tue vos mas	Vos seaux louent ce pot
Ce verre scie vos quais	Vos cas nient ce pot
Ce show tue vos mets	Vos quais tuent ce show
Ce rat noue vos bords	Vos cerfs noient ce chat
Ce pot noie vos mets	Vos mets lient ce pet
Ce pot loue vos mots	Vos cas lient ce show
Ce rat lie vos pas	Vos seaux voient ce corps
Ce verre voit vos mas	Vos seaux jouent ce rat
Ce bas loue vos bords	Vos cerfs jouent ce pet
Ce rat joue vos mots	Vos pas nouent ce pet
Ce verre loue vos cas	

References

- Aslin, R., Saffran, J., & Newport, L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, 9, 321–324.
- Bernal, S., Lidz, J., Millote, S., & Christophe, A. (2007). Syntax constrains the acquisition of verb meaning. *Language Learning and Development*, 3, 325–341.
- Boersma, P. & Weenink, D. (2008). Praat: Doing phonetics by computer (Version 5.0.25) (Computer program). <<http://www.praat.org/>>.
- Brown, R. (1957). Linguistic determinism and parts of speech. *Journal of Abnormal and Social Psychology*, 55, 1–5.
- Christophe, A., Millote, S., Bernal, S., & Lidz, J. (2008). Bootstrapping lexical and syntactic acquisition. *Language and Speech*, 51, 61–75.
- Cutler, A. (1993). Phonological cues to open- and closed-class words in the processing of spoken sentences. *Journal of Psycholinguistic Research*, 22, 109–131.
- Dabrowska, E. (2001). Learning a morphological system without a default: The Polish genitive. *Journal of Child Language*, 28, 545–574.
- Endress, A. D., & Mehler, J. (2009). The surprising power of statistical learning: When fragment knowledge leads to false memories of unheard words. *Journal of Memory and Language*, 60, 351–367.
- Garcia, J., Hankins, W. G., & Rusiniak, K. W. (1974). Behavioral regulation of the milieu interne in man and rat. *Science*, 185, 824–831.
- Gervain, J., Nespors, M., Mazuka, R., Horie, R., & Mehler, J. (2008). Bootstrapping word order in prelexical infants: A Japanese–Italian cross-linguistic study. *Cognitive Psychology*, 57, 56–74.
- Gillette, J., Gleitman, H., Gleitman, L. R., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, 73, 135–176.
- Gleitman, L. R. (1994). Words, words, words.... *Philosophical Transactions of the Royal Society of London, Series B*, 346, 71–77.
- Gleitman, L. R., & Gleitman, H. (1992). A picture is worth a thousand words, but that's the problem: The role of syntax in vocabulary acquisition. *Current Directions in Psychological Science*, 1, 31–35.
- Halberda, J. (2003). The development of a word-learning strategy. *Cognition*, 87, B23–B34.
- Hirsh-Pasek, K., Golinkoff, R. M., & Hollich, G. (2000). An emergentist coalition model for word learning: Mapping words to objects is a product of the interaction of multiple cues. In R. M. Golinkoff, K. Hirsh-Pasek, L. Bloom, L. Smith, A. Woodward, N. Akhtar, M. Tomasello, & G. Hollich (Eds.), *Becoming a word learner: A debate on lexical acquisition* (pp. 136–164). New York, NY: Oxford University Press.
- Johnson, E. K., & Jusczyk, P. W. (2001). Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language*, 44, 548–567.
- Johnson, E. K., & Seidl, A. H. (2009). At 11 months, prosody still outranks statistics. *Developmental Science*, 12, 131–141.
- Keen, R., & Machado, A. (1999). How pigeons discriminate the relative frequency of events. *Journal of the Experimental Analysis of Behavior*, 72, 151–175.
- Landau, B., & Gleitman, L. R. (1985). *Language and experience: Evidence from the blind child*. Cambridge, MA: Harvard University Press.
- Lidz, J., Gleitman, H., & Gleitman, L. R. (2003). Understanding how input matters: Verb learning and the footprint of universal grammar. *Cognition*, 87, 151–178.
- Liittschwager, J. C., & Markman, E. M. (1994). Sixteen and 24-month-olds' use of mutual exclusivity as a default assumption in second label learning. *Developmental Psychology*, 30, 955–968.
- Machado, A., & Cevik, M. (1997). The discrimination of relative frequency by pigeons. *Journal of the Experimental Analysis of Behavior*, 67, 11–41.
- Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14, 57–77.
- Markman, E. M., Wasow, J. L., & Hansen, M. B. (2003). Use of the mutual exclusivity assumption by young word learners. *Cognitive Psychology*, 47, 241–275.
- Naigles, L. G., & Kako, E. T. (1993). First contact in verb acquisition: Defining a role for syntax. *Child Development*, 64, 1665–1687.
- Nespor, M., & Vogel, I. (1986). *Prosodic phonology*. Dordrecht: Foris.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: MIT Press.
- Pinker, S. (1994). *The language instinct*. New York: William Morrow.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Savage, C., Lieven, E., Theakston, A., & Tomasello, M. (2003). Testing the abstractness of young children's linguistic representations: Lexical and structural priming of syntactic constructions? *Developmental Science*, 6–5, 557–567.
- Shi, R., Morgan, J. L., & Allopenna, P. (1998). Phonological and acoustic bases for earliest grammatical category assignment: A cross-linguistic perspective. *Journal of Child Language*, 25, 169–201.
- Shi, R., Werker, J., & Morgan, J. L. (1999). Newborn infants' sensitivity to perceptual cues to lexical and grammatical words. *Cognition*, 72, B11–B21.
- Shukla, M., Nespors, M., & Mehler, J. (2007). An interaction between prosody and statistics in the segmentation of fluent speech. *Cognitive Psychology*, 54, 1–32.
- Smith, L., & Yu, C. (2007). Infants rapidly learn words from noisy data via cross-situational statistics. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th annual cognitive science society* (pp. 653–658). Austin, TX: Cognitive Science Society.
- Stager, C. L., & Werker, J. F. (1997). Infants listen for more phonetic detail in speech perception than in word learning tasks. *Nature*, 388, 381–382.
- Swingle, D. (2007). Lexical exposure and word-form encoding in 1.5-year-olds. *Developmental Psychology*, 43, 454–464.
- Thiessen, E., & Saffran, J. (2003). When cues collide: Use of stress and statistical cues to word boundaries by 7- to 9-month-old infants. *Developmental Psychology*, 39, 706–716.
- Thothathiri, M., & Snedeker, M. (2008). Syntactic priming during language comprehension in three- and four-year-old children. *Journal of Memory and Language*, 58, 188–213.
- Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- Vickery, T. J., & Jiang, Y. V. (2009). Associative grouping: Perceptual grouping of shapes by association. *Attention, Perception & Psychophysics*, 71, 896–909.
- Voloumanos, A. (2008). Fine-grained sensitivity to statistical information in adult word learning. *Cognition*, 107, 729–742.
- Waxman, S. R., & Booth, A. E. (2001). Seeing pink elephants: Fourteen-month-olds' interpretations of novel nouns and adjectives. *Cognitive Psychology*, 43, 217–242.
- Waxman, S. R., & Booth, A. E. (2003). The origins and evolution of links between word learning and conceptual organization: New evidence from 11-month-olds. *Developmental Science*, 6, 128–135.
- Yang, C. D. (2004). Universal grammar, statistics or both? *Trends in Cognitive Sciences*, 8, 451–456.
- Yu, C., & Smith, L. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18, 414–420.